

Automatic 3D Reconstruction of Manifold Meshes via Delaunay Triangulation and Mesh Sweeping

Andrea Romanoni
Politecnico di Milano
Italy

andrea.romanoni@polimi.it

Amaël Delaunoy
ETH Zürich
Switzerland

amael.delaunoy@gmail.com

Marc Pollefeys
ETH Zürich
Switzerland

marc.pollefeys@inf.ethz.ch

Matteo Matteucci
Politecnico di Milano
Italy

matteo.matteucci@polimi.it

Abstract

In this paper we propose a new approach to incrementally initialize a manifold surface for automatic 3D reconstruction from images. More precisely we focus on the automatic initialization of a 3D mesh as close as possible to the final solution; indeed many approaches require a good initial solution for further refinement via multi-view stereo techniques. Our novel algorithm automatically estimates an initial manifold mesh for surface evolving multi-view stereo algorithms, where the manifold property needs to be enforced. It bootstraps from 3D points extracted via Structure from Motion, then iterates between a state-of-the-art manifold reconstruction step and a novel mesh sweeping algorithm that looks for new 3D points in the neighborhood of the reconstructed manifold to be added in the manifold reconstruction. The experimental results show quantitatively that the mesh sweeping improves the resolution and the accuracy of the manifold reconstruction, allowing a better convergence of state-of-the-art surface evolution multi-view stereo algorithms.

1. Introduction

Reconstructing the observed scene from a set of images is a widely studied problem in computer vision: it is known as *Structure from Motion* (SfM), when the aim is to reconstruct both the scene (the structure) and the camera poses (the motion), or as *multi-view stereo* (MVS), when camera poses are known. While SfM techniques provide a sparse reconstruction of the environment, MVS aims at reconstructing a dense and accurate model of the observed scene.

Since the datasets of [25] and [28] were made available, dense MVS has been faced with different approaches, some of which reaching very accurate results; however, some issues are still open, e.g., the initialization and the management of untextured regions and the handling of false

matches. According to the representation of the 3D model, MVS methods can be subdivided in *point-based*, *volumetric* and *mesh-based* methods. Point-based methods estimate the depth of each image pixel, reconstructing a point cloud [7, 30]; they provide accurate and dense point clouds, but the outcome is often redundant where the surface is flat, and they are not able to reconstruct untextured area. Volumetric methods, first proposed by [2], build the reconstruction by discretizing the space and labeling as matter or free space the voxels, as in [2], or the tetrahedra, as in [13]. Voxel-based methods require a huge amount of memory, and even if attempts to make the representation more compact exist [27], they are still not suitable for large-scale reconstruction. Tetrahedron-based methods are more compact, but the accuracy of the reconstruction still needs to be refined to reach state-of-the-art results. This latter approach is often used as an initialization to mesh-based methods in [33, 32, 24].

Mesh-based methods have been proven to be suitable to build continuous, high accurate reconstructions of both small objects and large-scale scenes [32, 33, 24]. These methods refine an existing mesh by minimizing an image similarity measure such as the Zero Mean Cross Correlation (ZNCC) [32, 20, 35] or the Sum of Squared Differences (SSD) [4, 3]. The initialization of the existing mesh is one of the major issues in state-of-the-art mesh-based MVS methods. For single object reconstruction, as in the Middlebury dataset [25], the initial mesh is usually estimated by the visual hull [14], but this is not applicable in more complex scene such as the ones in [28] or in large-scale scenarios.

One of the most effective and scalable mesh-based algorithm for multi-view stereo was proposed by Vu *et al.* [33]; in their work, the authors initialize the mesh evolution algorithm by extracting a very dense and noisy point cloud, then building a Delaunay triangulation out of it, and estimating the initial mesh via a *s-t* cut algorithm based on the work of [13]. This last work inspired other extensions such as [11]. The main problem with this approach is that the *s-t* cut does not guarantee the output mesh to be a manifold

while the *surface evolution* algorithm needs this property; indeed Vu *et al.* need to manually check the outcome of the *s-t* cut before surface evolution.

Some mesh-based algorithms, e.g., [19, 16], initialize the reconstruction with a point-based MVS such as CMVS [7] together with Poisson Reconstruction [12]. This process automatically provides an initialization usually close to be manifold, but still without guarantees, and it has issues related to redundancy of the points and non-scalability. In the literature, some methods to estimate a manifold mesh exist, but most of them rely on silhouettes [31, 6, 5], thus they are limited to small objects, or they are not scalable due to the usage of voxel-based reconstruction [10]. A suitable approach to estimate a manifold mesh has been recently proposed in [15, 17, 21, 22]: the authors incrementally reconstruct the scene from very sparse data which are the outcome of SfM algorithms. For many applications the reconstruction relying only on these points is sufficient, but for a surface evolution approach, it needs resolution and accuracy improvements. Indeed, to avoid getting stuck in local minima during the evolution process, also Vu *et al.* need to densify the point cloud extracted by the SfM algorithm.

In this paper we propose a novel and fully automatic approach to estimate a manifold, suitable for being a good initialization for a surface evolution mesh-based MVS algorithm. We bootstrap from a manifold reconstructed in a similar fashion to [17], and we refine it by iteratively sweeping the triangles of the manifold around their neighborhood looking for good stereo matches. This approach takes inspiration from the multi-directions plane-sweep of [8], where the authors sweep a set of planes choosing fixed directions related to the direction of the walls identified from Structure from Motion points; for each pixel and each plane they collect the matching costs among the views and choose the best stereo matching cost to estimate depth maps. Another interesting method that performs local plane sweeping is presented in [26], where, again, the output is a depth map. In the proposed approach we estimate directly a manifold mesh; the directions of the planes are automatically driven by the data, and we avoid to sweep on the whole scene, since we look for new 3D points in the neighborhood of the iteratively reconstructed manifold. Differently from [8], we directly recover a consistent dense 3D mesh instead of independent depth maps.

In Section 2 we summarize the reconstruction and refinement algorithm that builds incrementally a manifold and refines it. In Section 3 we show the results of our approach, underlying the improvement of the reconstruction accuracy thanks to the sweeping algorithm on both simulated and real datasets. In Section 4 we conclude pointing out possible future work.




Figure 1. The manifold property: the vertex V in (a) is regular ($ABCDEF$ is a closed path without cycles), while in (b) and (c) the vertex V is not regular since, in the former case the path $ABCDEFGHI$ is not closed, and in the latter $ABCDEFGHF$ has two cycles.

2. Proposed algorithm

In the proposed system, we iteratively alternate manifold reconstruction and mesh sweeping. First we describe how we estimate the initial mesh, then how we extract new points through mesh sweeping and how we add them, incrementally, to the manifold.

2.1. Manifold Reconstruction

Let recall that the manifold property holds if and only if the neighborhood of each surface point is homeomorphic to a disk. So a mesh, which is a discrete surface, is manifold if and only if each vertex v is *regular*, i.e., if and only if the edges opposite to v form a closed path without loops [17]. We show in Figure 1 (a) one example of a regular vertex and in Figure 1(b) and 1(c) two cases where the vertices are not regular and thus they break the manifold property.

In literature [9, 17], reconstruction from sparse data involves the estimation of a mesh which usually partitions the 3D Delaunay Triangulation of a set of sparse points between two sets: free space and matter. In this case, the sources of non-manifoldness can be induced by tetrahedra intersecting at a vertex, and by tetrahedra with a common edge (non-manifold edge). In both cases the non-manifoldness can be fixed, in principle, by cloning and offsetting the vertex of intersection, but two big drawbacks arise: this process would invalidate the Delaunay property, that we need to keep valid in an incremental setting, and, since many non-manifold vertices are generated, rearranging the triangulation and the visibility information for each non-manifold point becomes inefficient. For these two reasons we need a different approach to enforce the manifoldness.

We reconstruct a first manifold with the algorithm proposed in [17], slightly modified by a weighting scheme as in [22, 21], bootstrapping from the camera poses, the reconstructed points and the visibility estimated by a Structure from Motion algorithm¹. The manifold is reconstructed such that it partitions the 3D triangulation of the SfM points between the set O of *outside* tetrahedra, i.e., the subset of the free space outside the manifold (not all the free space tetrahedra will be part of the space outside the manifold),

¹In our implementation we use VisualSFM [34].




Figure 2. Example of manifold sweeping with respect to the camera C

and the complementary set I of inside tetrahedra (i.e., the remaining tetrahedra that represent the matter together with the free space tetrahedra which would invalidate the manifold property). A weight is associated to each tetrahedron and it keeps track of the visibility information, i.e., the camera-to-point viewing rays; in the following, a tetrahedron belongs to free space if its weight is higher than a threshold t_w (in our case $t_w = 4.0$).

The manifold is obtained by three steps. (1) *Point Insertion*: build the 3D Delaunay triangulation of the 3D SfM points. (2) *Ray tracing and tetrahedra weighting*: for each camera-to-point viewing ray, add a weight $w_1 = 4.0$ to the intersected tetrahedra and a weight $w_2 = 0.5$ to their neighboring tetrahedra. Such weighting scheme acts as a smoother of the visibility and avoids the creation of visual artifacts (see [18] for a detailed discussion about visual artifacts). (3) *Growing*: initialize a queue Q starting from the tetrahedron with the higher weight. Then: (a) pick the tetrahedron with highest weight from Q and add it to O only if the resulting surface between O and I remains manifold; (b) add the neighboring tetrahedra to the queue Q , otherwise discard it; (c) continue iteratively until Q is empty.

2.2. Mesh Sweeping

After each manifold reconstruction step, we apply the novel mesh sweeping algorithm to look for new 3D points in the neighborhood of the manifold mesh, which we assume to be close to the true model of the scene. These points will be added to the reconstruction at the next manifold reconstruction iteration.

For each camera, we sweep the visible part of mesh \mathcal{M} along the viewpoint direction (see Figure 2). Given a camera located in C , we sweep each visible facet f of \mathcal{M} by an amount multiple of α (in our case $\alpha = 3\text{cm}$). Let v^1, v^2 and v^3 be the vertices of f , and let θ^i be the angle between the normal n of the facet and the ray d^i from the camera center to the vertex v^i . The new vertex v_j^i of the swept facet is now




Figure 3. 3D points extraction process after the mesh sweeping.

computed as: $v_j^i = v^i + \alpha_j \cdot \cos(\theta^i) \cdot d^i$, where $\alpha_j = \alpha \cdot k_j$, with $k_j \in \mathbb{N}$ and $-\frac{N}{2} < k_j < \frac{N}{2}$, is the distance swept (let note that is fixed for all the facets). In our case we fixed $N = 20$.

Since all the visible vertices of each triangle are swept by the same amount along the camera-to-point direction, no self-intersections among facets has been induced.

The sweeping process produces a set of $N + 1$ meshes $\mathbb{M} = \{\mathcal{M}_1, \dots, \mathcal{M}_j, \dots, \mathcal{M}_{N+1}\}$. For each mesh \mathcal{M}_j we perform pairwise stereo matching between each camera C and the two nearest cameras. Let I_C be the image seen by C , and I_k the image seen by the camera C_k , that is one of the two cameras nearest to C , therefore $k \in \{1, 2\}$. For each camera k and mesh j , we compute I_k^j as the reprojection of image I_k through the mesh \mathcal{M}_j into camera C . Then, we compute the Normalized Cross Correlation (NCC) image I_C and the reprojection I_k^j . The NCC is computed pixel-by-pixel and is weighted by a Gaussian kernel as in [20] (the σ for the Gaussian kernel is $\sigma = 8\text{px}$). The last step aims at collecting the new 3D points. We create a $n_r \times n_c$ grid over the N images of pixel-by-pixel NCCs computed for each mesh \mathcal{M}_j . We collect in each tile the image point with the best NCC above a threshold $t_{\text{NCC}} = 0.98$ among all the N images; therefore, we obtain at most $n_r \times n_c$ points. For each collected point, we retrieve its 3D position from the corresponding mesh that generated it.

We are able to perform the whole process image-based, i.e., we act on the image and not on the mesh, so that the most complex computations become independent from the size of the mesh, and they are self-adaptive with respect to the image resolution. This also allows our approach to be adaptive in term of mesh resolution output.

Mesh sweeping requires some parameters. N and α need to be fixed in order to keep the sweeping local and, at the same time, to induce an appreciable sweeping of the mesh w.r.t. the image resolution. The values of N or α are also related to the observed scene and the image resolution, but the sweeping algorithm is quite robust against these two pa-




Figure 4. Simulated dataset results: (a) textured image, (b) ground-truth, (c) initial mesh, (d) after mesh sweeping, (e) after photometric refinement.

parameters. On the other hand, σ , t_{NCC} , n_r and n_c concur to determine the trade off between precision and recall of the newly estimated 3D points. The threshold t_{NCC} needs to be high so to avoid outliers, even if a good choice of the tile sizes, e.g., 100x100px, manages to make the algorithm robust to outliers and to small variations of parameter setting. Finally, the value of σ is usually related to the image resolution, but in our experiments, with different resolutions, we found that the same value of 8px is a good compromise between cross-correlation sensitivity and smoothness.

2.3. Incremental refinement of the manifold

To refine the initial mesh, we iterate between the mesh sweeping algorithm, which outputs a set P of new 3D points, and the incremental reconstruction (described in this section) until convergence is reached (i.e., no more points are available to be added or the number of iteration is greater than $it_{max} = 15$).

The insertion of a point $p \in P$ into the Delaunay triangulation causes the removal of a set D of tetrahedra breaking the Delaunay property. The surface between $O \setminus D$ and $I \cup D$ is not guaranteed to be manifold anymore. In order to avoid this, as the authors in [17], we define a list of tetrahedra $E \supset D$ and apply the *Shrinking* procedure, i.e., the inverse of Growing: we subtract iteratively from O only the tetrahedra $\Delta \in E$ keeping the manifoldness valid. After this process, it is likely that $D \cap O = \emptyset$. Whenever $D \cap O \neq \emptyset$ the point p is not added to the triangulation and it is discarded. Once all points in P have been processed, the queue Q is initialized with the tetrahedra $\Delta \in T \setminus O$ such that $\Delta \cap \delta O \neq \emptyset$, and the Growing process starts as explained in Section 2.1.

3. Experimental evaluation

The proposed algorithm provides an automatic initialization for mesh-based dense MVS algorithms. We implemented our approach in C++ using the CGAL library [29] for all the computations involving the Delaunay Triangulation and manifold extraction. This allows us to exploit efficient data structures and algorithms for the initial manifold reconstruction tasks. Conversely, the proposed mesh sweeping algorithm exploits the power of GPU computing; we implemented the main steps with the GLSL shading language of OpenGL [23]: a geometry shader computes the normal; the vertex shaders sweep the mesh along the camera viewing rays; the fragment shaders implement the re-projection from one image to another through the mesh surface, the NCC computation and the thresholding.

We performed experiments on both synthetic and real datasets on a 4 Core i7-920 CPU at 2.6Ghz (8M Cache), with 24GB of DDR3 SDRAM and NVIDIA TitanX. Synthetic experiments aim at showing the effectiveness of the mesh sweeping approach to refine a rough manifold. The experiments on real data show the scalability of the proposed method and its effectiveness compared to CMVS [7], used as automatic initialization in mesh-based algorithms [19, 16] together with Poisson reconstruction [12].

In Figure 4 we provide quantitative evaluation for the two synthetic datasets. The ground-truth meshes we used to generate the image are two pyramids without the squared base facet (Figure 4(b)), pointing downward for the first dataset, and upward in the second one; the dimensions of these pyramid are 2x2m of squared base and 0.3m height. The initial point cloud in both cases is made up by the four vertices of the base. We bootstrap from only these four points, to show the impact of the proposed approach even




Figure 5. Cumulative distributions of error in the synthetic datasets.




Figure 6. Cumulative distribution of errors in the fountain-P11 dataset.

if few sparse points are available. In both cases the initial manifold extracted is a square corresponding to the base of the pyramid (Figure 4(c)). The final reconstructions are shown in Figure 4(d) even if some little artifacts are created due to the noise in the matching process, the extracted surface is close to the ground-truth. We computed the distance of the meshes with respect to the ground truth by comparing the depth of each pixel from the first camera, similarly to the comparison method proposed in [28]. In Table 1 we compare the accuracy of the initial mesh with the accuracy after the proposed mesh sweeping: our method improves significantly the initial mesh; this result is also supported by the improvement on the error distributions in Figure 5, where the errors are reported as a percentage of depths such that their error e w.r.t. the ground truth is $e < m * \text{Sigma}$ (Sigma = 0.06m). In Table 2 we applied the photometric refinement of [33] to the two meshes and we show the results: the convergence of the refinement after the mesh sweeping is faster and the the reconstruction more accurate.

Table 1. Results of the proposed mesh sweeping in the synthetic datasets. All errors are in meters.

	Downward Pyramid			Upward Pyramid		
	MEA	MRE	RMS	MEA	MRE	RMS
w/o mesh sweeping	0.108	0.072	0.131	0.135	0.114	0.159
w/ mesh sweeping	0.013	0.008	0.025	0.028	0.022	0.049
improvement	88 %	89 %	81 %	79 %	81 %	69 %

Table 2. Photometric refinement with or without the mesh sweeping in the synthetic datasets. All errors are in meters.

Initial mesh	n. iter	Downward Pyramid			Upward Pyramid		
		MEA	MRE	RMS	MEA	MRE	RMS
w/o mesh sweeping	500	0.089	0.060	0.123	0.124	0.106	0.161
w/o mesh sweeping	5000	0.013	0.008	0.025	0.019	0.010	0.060
w/ mesh sweeping	500	0.002	0.001	0.008	0.012	0.009	0.029

Table 3. Comparison between CMVS+Poisson, initial mesh and final mesh. Errors are expressed in Mean Absolute Error (MAE), Mean Relative Error (MRE) and Root Mean Square error (RMS).


	MEA (m)	MRE (m)	RMS (m)	num. vertices
CMVS [7] + Poisson	0.160	0.019	0.317	253410
initial mesh [22]	0.163	0.096	0.330	5880
final mesh (Proposed)	0.094	0.012	0.210	22336

Table 4. Photometric refinement with or without the mesh sweeping in the fountain dataset. All errors are in meters.

Initial mesh	n. iter.	MEA (m)	MRE (m)	RMS (m)	num. vertices
w/o mesh sweeping	500	0.130	0.157	0.304	1M
w/ mesh sweeping	500	0.069	0.008	0.196	1M

We tested the proposed approach on the *fountain-P11* dataset provided in [28]: the ground-truth is available and the evaluation is performed by comparing the depth map generated from the 6th point of view as suggested in [28], such that the results of the comparison are independent from the quality of the camera calibration (a detailed discussion is available in [28]). We initialized the manifold reconstruction with the sparse point cloud data estimated with VisualSFM [34].

Table 3 shows the accuracy of the initial manifold, estimated in a similar fashion to [22], the manifold after the mesh sweeping and the reconstruction performed by CMVS+Poisson, where the octree depth of the Poisson reconstruction is 11 (see [12]). In Figure 6 we report the cumulative distribution of the errors for each algorithm with the same convention as before. As Figure 7 shows, the initial manifold represents accurately the wall, while some parts of the fountain present artifacts, indeed, the accuracy of this initial guess is lower than the outcome of CMVS+Poisson. The mesh sweeping refinement improves the mesh accuracy, both respect to the initial mesh and to CMVS+Poisson reconstruction, even if the resolution of the latter is significantly higher. Our approach avoids to reconstruct redundant data, differently from the point-based CMVS, which reconstruct as much points as possible, even if they are coplanar. Table 4 shows that the proposed approach improves the convergence of the variational photo-



CMVS Poisson Initial manifold After Mesh Sweeping Photometric Refinement

Figure 7. Reconstruction results. First row: depth errors darker blue encodes lower errors. Second row: untextured reconstruction. Third row: textured reconstruction.

metric refinement described in [33] leading to a more accurate reconstruction. The mean timing per-iteration is 318s for mesh sweeping and 55s for manifold reconstruction: the mesh sweeping is the most demanding task, but thanks to its GPU implementation it would easily take advantage from more powerful hardware.

Finally we tested the scalability of the proposed approach with a real dataset, named Dagstuhl, of 68 1600x1200 images (see Figure 8). We tested the algorithm considering different subsets of frames to understand experimentally the dependence with the number of cameras of the proposed algorithm and the two steps (manifold reconstruction and mesh sweeping) separately. As Figure 9 points out, both steps and the whole algorithm show experimentally a linear relation with the number of cameras; this suggests that our approach is scalable and suitable for large-scene reconstructions.

4. Future Works and Conclusion

In this paper we proposed a new approach to reconstruct an initialization for mesh-based multi-view stereo algorithm such that the whole pipeline is fully automatic, taking into account that state-of-the-art algorithms are not able to evolve correctly non-manifold meshes. We build a first rough mesh with an incremental manifold reconstruction algorithm, then, we iterate between the novel mesh sweeping refinement and the manifold reconstruction. The former sweeps the mesh along the cameras viewing direction and seeks new pairwise stereo matches in these swept meshes; the corresponding 3D points and visibility information are added to the incremental manifold surface algorithm. Our approach improves the initial mesh accuracy and resolution, and also with respect to the reconstruction based on



Figure 8. One frame of the Dagstuhl and reconstruction with the proposed method.




Figure 9. Processing time for the Dagstuhl dataset.

CMVS+Poisson, usually adopted as mesh-based MVS initialization. The presented algorithm has the strong advantage of being suitable for large scale, resolution-adaptive, and incremental 3D reconstruction scenarios, while keeping the manifold properties.

As future work we would improve the robustness with respect to untextured regions, in particular when more than one main untextured plane is present, and we would like to test other stereo matching measures. Moreover we would like to test in our scenarios some ideas from [1], such as the addition of the ghost primitives.

5. Acknowledgements

We thank NVIDIA who kindly supported this work through the Hardware Grant Program.

References

- [1] A.-L. Chauve, P. Labatut, and J.-P. Pons. Robust piecewise-planar 3d reconstruction and completion from large-scale unstructured point data. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1261–1268. IEEE, 2010. 6
- [2] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312. ACM, 1996. 1
- [3] A. Delaunoy and E. Prados. Gradient flows for optimizing triangular mesh-based surfaces: Applications to 3d reconstruction problems dealing with visibility. *International journal of computer vision*, 95(2):100–123, 2011. 1
- [4] A. Delaunoy, E. Prados, P. G. I. Piracés, J.-P. Pons, and P. Sturm. Minimizing the multi-view stereo reprojection error for triangular surface meshes. In *BMVC 2008-British Machine Vision Conference*, pages 1–10. BMVA, 2008. 1
- [5] C. H. Esteban and F. Schmitt. Silhouette and stereo fusion for 3d object modeling. *Computer Vision and Image Understanding*, 96(3):367–392, 2004. 2
- [6] Y. Furukawa and J. Ponce. Carved visual hulls for image-based modeling. In *Computer Vision—ECCV 2006*, pages 564–577. Springer, 2006. 2
- [7] Y. Furukawa and J. Ponce. Accurate, dense, and robust multiview stereopsis. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(8):1362–1376, 2010. 1, 2, 4, 5
- [8] D. Gallup, J.-M. Frahm, P. Mordohai, Q. Yang, and M. Pollefeys. Real-time plane-sweeping stereo with multiple sweeping directions. In *Computer Vision and Pattern Recognition, 2007. CVPR’07. IEEE Conference on*, pages 1–8. IEEE, 2007. 2
- [9] C. Hoppe, M. Klopschitz, M. Donoser, and H. Bischof. Incremental surface extraction from sparse structure-from-motion point clouds. *Proc. BMVC*, 2013. 2
- [10] A. Hornung and L. Kobbelt. Hierarchical volumetric multi-view stereo reconstruction of manifold surfaces based on dual graph embedding. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 1, pages 503–510. IEEE, 2006. 2
- [11] M. Jancosek and T. Pajdla. Multi-view reconstruction preserving weakly-supported surfaces. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 3121–3128. IEEE, 2011. 1
- [12] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, volume 7, 2006. 2, 4, 5
- [13] P. Labatut, J.-P. Pons, and R. Keriven. Efficient multi-view reconstruction of large-scale scenes using interest points, delaunay triangulation and graph cuts. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pages 1–8. IEEE, 2007. 1
- [14] A. Laurentini. The visual hull concept for silhouette-based image understanding. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 16(2):150–162, 1994. 1
- [15] M. Lhuillier and S. Yu. Manifold surface reconstruction of an environment from sparse structure-from-motion data. *Computer Vision and Image Understanding*, 117(11):1628–1644, 2013. 2
- [16] Z. Li, K. Wang, W. Zuo, D. Meng, and L. Zhang. Detail-preserving and content-aware variational multi-view stereo reconstruction. *arXiv preprint arXiv:1505.00389*, 2015. 2, 4
- [17] V. Litvinov and M. Lhuillier. Incremental solid modeling from sparse and omnidirectional structure-from-motion data. In *BMVC*, 2013. 2, 4
- [18] V. Litvinov and M. Lhuillier. Incremental solid modeling from sparse structure-from-motion data with improved visual artifacts removal. In *International Conference on Pattern Recognition (ICPR)*, 2014. 3
- [19] R. Pan and G. Taubin. Automatic segmentation of point clouds from multi-view reconstruction using graph-cut. *The Visual Computer*, pages 1–9, 2015. 2, 4
- [20] J.-P. Pons, R. Keriven, and O. Faugeras. Multi-view stereo reconstruction and scene flow estimation with a global image-based matching score. *International Journal of Computer Vision*, 72(2):179–193, 2007. 1, 3
- [21] A. Romanoni and M. Matteucci. Efficient moving point handling for incremental 3d manifold reconstruction. In *Image Analysis and Processing—ICIAP 2015*, pages 489–499. Springer, 2015. 2
- [22] A. Romanoni and M. Matteucci. Incremental reconstruction of urban environments by edge-points delaunay triangulation. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pages 4473–4479. IEEE, 2015. 2, 5
- [23] R. J. Rost, B. M. Lincea-Kane, D. Ginsburg, J. M. Kessenich, B. Lichtenbelt, H. Malan, and M. Weiblen. *OpenGL shading language*. Pearson Education, 2009. 4
- [24] N. Salman and M. Yvinec. Surface reconstruction from multi-view stereo of large-scale outdoor scenes. *International Journal of Virtual Reality*, 9(1):19–26, 2010. 1
- [25] S. M. Seitz, B. Curless, J. Diebel, D. Scharstein, and R. Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In *Computer vision and pattern recognition, 2006 IEEE Computer Society Conference on*, volume 1, pages 519–528. IEEE, 2006. 1
- [26] S. N. Sinha, D. Scharstein, and R. Szeliski. Efficient high-resolution stereo matching using local plane sweeps. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 1582–1589. IEEE, 2014. 2
- [27] F. Steinbrucker, J. Sturm, and D. Cremers. Volumetric 3d mapping in real-time on a cpu. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 2021–2028. IEEE, 2014. 1
- [28] C. Strecha, W. von Hansen, L. Van Gool, P. Fua, and U. Thoennessen. On benchmarking camera calibration and multi-view stereo for high resolution imagery. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008. 1, 5
- [29] The CGAL Project. *CGAL User and Reference Manual*. CGAL Editorial Board, 4.5 edition, 2014. 4
- [30] E. Tola, C. Strecha, and P. Fua. Efficient large-scale multi-view stereo for ultra high-resolution image sets. *Machine Vision and Applications*, 23(5):903–920, 2012. 1

- [31] G. Vogiatzis, P. H. Torr, and R. Cipolla. Multi-view stereo via volumetric graph-cuts. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 2, pages 391–398. IEEE, 2005. 2
- [32] H. H. Vu, R. Keriven, P. Labatut, and J.-P. Pons. Towards high-resolution large-scale multi-view stereo. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 1430–1437. IEEE, 2009. 1
- [33] H. H. Vu, P. Labatut, J.-P. Pons, and R. Keriven. High accuracy and visibility-consistent dense multiview stereo. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(5):889–901, 2012. 1, 5, 6
- [34] C. Wu. Visualsfm: A visual structure from motion system, 2011. 2, 5
- [35] A. Zaharescu, E. Boyer, and R. Horaud. Transformesh: a topology-adaptive mesh-based approach to surface evolution. In *Computer Vision—ACCV 2007*, pages 166–175. Springer, 2007. 1